



Commitment in the household: Evidence from the effect of inheritances on the labor supply of older married couples[☆]

David M. Blau^{a,*}, Ryan M. Goodstein^b

^a Department of Economics, Ohio State University, United States

^b Federal Deposit Insurance Corporation, United States

HIGHLIGHTS

- Receiving an inheritance causes a reduction in the recipient's labor supply.
- There is no impact of the inheritance on the labor supply of the recipient's spouse.
- This is evidence against the feasibility of full commitment within the household.
- The results are consistent with a model of limited commitment.

ARTICLE INFO

Article history:

Received 17 September 2015

Received in revised form 24 May 2016

Accepted 16 August 2016

Available online 24 August 2016

Keywords:

Retirement

Commitment

Inheritance

Household bargaining

Labor force participation

ABSTRACT

We study the effect of receiving an inheritance on the labor force participation (LFP) of both the recipient and the recipient's spouse in a population of older married couples. An inheritance is not subject to laws in the U.S. governing division of marital property at divorce, because it is not acquired with income earned during marriage. Hence it plays the role of a "distribution factor" in the intrahousehold allocation of resources, increasing bargaining power of the recipient. Controlling for inheritance expectations, we interpret the receipt of an inheritance as a shock to wealth. Our results indicate that receiving an inheritance reduces LFP of the recipient by four percentage points, comparable in magnitude to the effect of a self-reported decline in health. However, an inheritance has little or no effect on LFP of the spouse. These estimates are inconsistent with a dynamic, collective model of the household in which spouses have the ability to commit to an ex ante efficient allocation. The results are consistent with a model of limited commitment in which a shock to household resources can alter bargaining power. We discuss the implications for reform of Social Security spouse and survivor benefits.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Cooperative bargaining models of intrahousehold resource allocation have been applied with increasing frequency to analyze and interpret intertemporal behavior of households in an environment of uncertainty (see [Browning et al., 2014](#), for an overview). A key issue in this setting is whether household members are able to fully commit to a resource allocation plan (a "contract") agreed upon at the time the household is formed. If spouses can commit to a state-contingent resource allocation plan, then their relative bargaining power at the time of marriage determines the effects of subsequent income and other shocks on intrahousehold allocations. Such shocks would have

wealth and/or substitution effects, but they would not cause renegotiation of the original contract.¹

Commitment is an important issue because, as [Mazzocco \(2007\)](#) points out, it determines the impact of public policies that shift control of resources within the household. If households are able to commit to an ex ante efficient resource allocation plan, then policies that intentionally or unintentionally change control of resources within the household will have limited impact on intrahousehold resource allocation.² However, [Voena \(2015\)](#) argues that unilateral divorce laws, which are ubiquitous in the US today, limit the ability

[☆] The views expressed in this paper are those of the authors and do not necessarily reflect the views of the FDIC.

* Corresponding author at: Ohio State University, Department of Economics, 445 Arps Hall, 1945 N. High St., 43210-1172 Columbus, OH, United States.

E-mail addresses: blau.12@osu.edu (D.M. Blau), rgoodstein@fdic.org (R.M. Goodstein).

¹ See [Marcet and Marimon \(2011\)](#) for a general discussion of contracting problems in which agents are subject to intertemporal participation or other constraints that affect the set of feasible contracts. Of course, a contract can always be renegotiated by mutual consent, regardless of commitment ability.

² Such policies also operate via the budget constraint, so they will have wealth and/or substitution effects. And they will affect the initial distribution of bargaining power in households formed after implementation of the new policy.

of spouses to commit. In this legal environment, a shock that increases the relative value of the outside alternative for one spouse may result in a binding participation constraint, causing a shift in bargaining power within the household. In a cooperative bargaining framework this will cause renegotiation of the contract, leading to an ex post efficient outcome, given the new distribution of bargaining power. The new outcome could involve divorce, if that is efficient, or a reallocation of decision power toward the spouse whose participation constraint binds. But the inability to commit to an efficient resource allocation plan will lead to an ex ante inefficient outcome. For example, specialization of one spouse in home production activities and the other in the labor market may be optimal, but if the spouse who specializes in the market cannot commit to remaining in the household when his earnings are high, the optimal degree of specialization will not occur.³

Previous empirical studies of intertemporal household behavior in the cooperative bargaining framework have either assumed that spouses have full commitment ability and imposed the assumption in a structural estimation approach (Casanova, 2010; van der Klaauw and Wolpin, 2008), or have tested for full commitment by analyzing the implications for consumption or time allocation Euler equations (Lich-Tyler, undated; Lise and Yamada, 2014; Mazzocco, 2007).⁴ The drawback of the first approach is clear: if full commitment is not feasible, the model is misspecified. A drawback of the second approach is that Euler equation methods are not well-suited to analyze labor supply. Labor supply decisions are typically discrete, especially at older ages, where the most common pattern of retirement is abrupt and complete withdrawal from the labor force.

Our paper introduces a new approach to empirical analysis and testing of commitment in married-couple households. We estimate the impact of receiving an inheritance on the labor force participation (LFP) decisions of older individuals and their spouses. Inheritances provide a useful new source of identification for studying commitment, because they are not subject to marital property law in the US. In most US states these laws specify that earnings during marriage and the assets acquired with those earnings are community property, divided equally or “equitably” between the spouses in the event of divorce, regardless of which spouse formally holds title to the asset (Mazzocco, 2007; Voena, 2015). For example, an employer-provided pension account held by one spouse is considered community property in the event of a divorce if the job was held during the marriage. In contrast, inheritances belong exclusively to the recipient since they were not acquired with earnings during marriage. Inheritances unambiguously increase the value of the outside option of the recipient but not of the spouse. Given the exclusion of inheritances from laws governing marital property in the US, inheritances are not contractible. This implies an inability to commit, at least with respect to inheritances.

³ The legal environment governing household dissolution and property division for cohabiting couples is very different than for married couples. Hence we do not analyze or discuss cohabiting couples, although many of the same issues are relevant.

⁴ An exception is Lundberg et al. (2003), who analyze the change in household consumption expenditure following retirement of the husband, and interpret the results in terms of an intertemporal bargaining model without the assumption of commitment. Our approach is similar, as it develops a test based on a model and imposes minimal assumptions in the estimation. Mazzocco et al. (2007) estimate a dynamic collective labor supply model without commitment for young couples. Gemici (2011) estimates a dynamic cooperative Nash bargaining model of family labor supply and migration. She assumes that utility is transferable, leading to an efficient outcome despite lack of commitment ability. Voena (2015) specifies a model in which commitment is assumed to be feasible in a mutual-consent divorce regime and infeasible in a unilateral divorce regime. Several papers have used a non-cooperative bargaining approach to modeling retirement behavior of couples: e.g. Gallipoli and Turner (2013) and Gustman and Steinmeier (2009). By construction, there is no commitment ability in such models.

We use inheritances to test for commitment in a discrete choice labor supply framework. Our approach is similar in spirit to Mazzocco (2007), but our test is for labor supply rather than consumption, imposes weaker assumptions, and uses a new source of identification.⁵ Under the null hypothesis of full commitment ability, the effect on the husband's LFP of an unexpected inheritance received by him should be equal to the effect on his LFP of an unexpected inheritance received by his wife, and conversely for the wife's LFP. Under full commitment, decision power at the time of marriage determines the allocation of resources in the couple's state-contingent contract. For example, if both spouses perceive a high probability that the wife will inherit a large sum in the future, her decision power at the time of the marriage will tend to be relatively high. The actual receipt of an inheritance will affect LFP of the spouses via wealth effects as determined by their initial decision power, regardless of which spouse is the recipient. A pattern in which a husband's inheritance affects only his LFP and a wife's inheritance affects only her LFP is inconsistent with full commitment, but is consistent with a limited commitment model in which contracts are renegotiated when a shock causes a participation constraint to bind. We develop a simple model in the next section to illustrate this point.

Our empirical analysis uses longitudinal data from the Health and Retirement Study (HRS) on inheritances and inheritance expectations of both spouses in married-couple households. Controlling for inheritance expectations, we interpret inheritance receipt as a shock. This is a rare example of a measureable household resource shock that unambiguously accrues to a specific household member. It is important to focus on inheritance shocks, because an inheritance that is anticipated at the beginning of the marriage should not affect the bargaining power of the recipient at the time of receipt.

We find that receiving an inheritance reduces LFP of the recipient by 4 percentage points, and has virtually zero impact on LFP of the spouse, controlling for inheritance expectations, lagged LFP, lagged inheritances, household wealth, and many other determinants of labor supply. The estimates of the own-inheritance effects for husbands and wives are similar in magnitude. We reject the null hypothesis of full commitment in many though not all specifications. The results are quite robust to alternative definitions of employment, alternative regression specifications, and alternative estimation approaches.

This finding confirms results from previous studies that have analyzed the impact of changes in control over resources within the household resulting from exogenous policy changes, but our context is quite different. Previous studies have focused mainly on spending on children as a function of who controls income entering the household.⁶ Our study is one of the first to focus on the impact of control over household resources on LFP.⁷ We contribute to the literature on commitment by using unanticipated inheritances as a new source of identifying information, and by studying retirement, a major life decision. In the concluding section we discuss reform of Social Security spouse and survivor benefits as an important example of a policy change the effects of which depend on commitment ability.

Our paper is most closely related to two recent papers. Brown et al. (2010) exploit the HRS survey data on anticipated and actual receipt

⁵ The assumptions of the Euler equation approach include intertemporal separability of preferences and the absence of liquidity constraints.

⁶ See Lundberg et al. (1997), Bobonis (2009), Dufo (2003), and Dufo and Udry (2004).

⁷ Chiappori et al. (2002) use a static framework to analyze the effects of various “distribution factors” on hours of work in two-earner households, but they do not study the participation decision. A number of studies treat the ratio of the spouse's wage rates as a distribution factor, but the wage ratio is unlikely to be exogenous. Lise and Yamada (2014) study commitment in a model of time allocation, using deviations of wage growth from the path anticipated at the time of marriage as a measure of resource shocks. To implement this approach, they specify a wage forecasting model that is assumed to be used by individuals. An advantage of our approach is that we do not have to make assumptions about how expectations are formed.

of inheritances to construct a measure of unanticipated inheritances. They find that receipt of an unanticipated inheritance leads to an increase in labor force exit at older ages.⁸ They focused on the effect of *household-level* receipt of an inheritance, so their estimate represents the average effect of inheritance receipt on household labor supply. We extend their analysis to estimate both the own and cross-spouse effects of inheritance shocks, disaggregated by the identity of the recipient. Cesarini et al. (2015) use administrative data on Swedish Lottery winners between the ages of 21 and 64. Their data allow them to estimate the impact on both the winner and the winner's spouse. They find a negative effect of winning the lottery on the spouse's labor earnings, but the effect is smaller than the effect on the lottery winner's earnings. This finding is similar to our results described above, and confirms the absence of full commitment ability.

2. Model

We develop a stylized model to motivate our test of commitment.⁹ Consider a two-person household and a two-period horizon. Spouse i has a period utility function defined over consumption and hours of work: $u_i(c_{it}, h_{it})$. For simplicity, there are no household public goods. We focus on the hours of work choice, but the extension to the discrete work decision is straightforward. We assume cooperative behavior that leads to a Pareto efficient outcome (see Browning et al., 2014, for a survey of this literature). This implies a formulation in which the spouses choose consumption and hours of work each period to maximize a weighted sum of the spouses' expected present discounted values of lifetime utility,

$$\mu_1 E \sum_{t=1}^2 u_1(c_{1t}, h_{1t}) + \mu_2 E \sum_{t=1}^2 u_2(c_{2t}, h_{2t})$$

with respect to consumption and hours of work each period, subject to constraints specified below. μ_i is the ex ante bargaining power or Pareto weight of person i at the time the match is formed, which is a function of distribution factors to be specified below, and E is the expectations operator. For simplicity we have assumed no discounting, and we will also assume that the interest rate is zero.

Resources are derived from a household-level endowment A_0 , earnings $w_{it}h_{it}$, and spouse-specific inheritances I_i . The household faces no liquidity constraint, but must be solvent at the end of period 2. Assume for simplicity that wage rates are constant over time. Inheritances are random variables realized at the beginning of period 2, before period-2 choices are made. The joint probability density function (pdf) of inheritances is $f(I_1, I_2)$. Inheritances are the only source of uncertainty in the model. The budget constraint in the first period is

$$c_{11} + c_{21} = A_0 + w_1 h_{11} + w_2 h_{21} - A_1,$$

where A_1 is savings, and the state-contingent budget constraint in period 2 is

$$c_{12} + c_{22} = A_1 + w_1 h_{21} + w_2 h_{22} + I_1^* + I_2^*,$$

where I_i^* is the realization of the random variable I_i . These constraints are based on the assumption that resources are pooled within the household, a key element of cooperative bargaining models.

Following Chiappori et al. (2002), define a distribution factor as a variable that affects the intrahousehold decision process but does not influence preferences or the couple's joint budget constraint. The key distribution factor in this model is f . The greater the likelihood that spouse i will receive an inheritance, as measured by the joint pdf, the greater is her ex ante bargaining power at the time the marriage is formed.¹⁰ The assumption of commitment means that the Pareto weights $\mu_i(f)$ are fixed at the time of marriage: the couple commits to an allocation plan, and the realization of the inheritance outcome does not cause renegotiation. With this assumption, the model is complete and can be solved recursively. The household's problem in period 2 is

$$\max \mu_1(f) u_1(c_{12}, h_{12}) + \mu_2(f) u_2(c_{22}, h_{22})$$

$$\{c_{i2}, h_{i2}\}_{i=1,2}$$

subject to the period 2 budget constraint.

The solution can be written in the form of state-contingent value functions at the beginning of period 2, after realization of inheritances, $V_{i2}(\Phi, I_1^*, I_2^*)$, where Φ is the vector of state variables known at the end of period 1: $\Phi = \{A_1, w_1, w_2\}$. The key empirical implication of commitment derives from the fact that inheritance *realizations* enter the problem only through the period-2 budget constraint, where they appear additively. The ex ante probability distribution of inheritances affects bargaining power, but under commitment the realizations do not. The realizations have wealth effects only. Hence in the case of commitment we can rewrite $V_{i2}(\Phi, I_1^*, I_2^*)$ as $V_{i2}^*(\Phi, I_1^* + I_2^*)$, indicating that the recipient of the inheritance is irrelevant.

If commitment is not possible, there are participation constraints in period 2: $u_i(c_{i2}, h_{i2}) \geq u_{i2}^*(I_i^*)$, $i = 1, 2$, where u_{i2}^* is the level of utility associated with the outside option of spouse i . We show the dependence of the utility of the outside option on the inheritance realization to emphasize the point that receiving an inheritance increases the value of the outside option. The key consideration in the absence of commitment is whether one of the spouses receives an inheritance shock large enough to cause a participation constraint to bind. If neither spouse experiences this event, the solution is identical to the commitment case. If both spouses receive such an inheritance, then both spouses prefer the outside option. If the outside option is divorce, the marriage ends and we don't observe the household in the data in period 2. Thus we focus here on the case in which one and only one spouse receives such an inheritance.¹¹ Let λ_i be the multiplier on the participation constraint for spouse i . Following Mazzocco (2007) and Marimon and Marimon (2011), the optimization problem in period 2 in the absence of commitment ability can be written as

$$\max_{\{c_{i2}, h_{i2}\}} \sum_{i=1}^2 \mu_i u_i(c_{i2}, h_{i2}) + \lambda_i (u_i(c_{i2}, h_{i2}) - u_{i2}^*)$$

subject to the budget constraint. If a participation constraint is not binding ($u_i(c_{i2}, h_{i2}) - u_{i2}^* > 0$), then $\lambda_i = 0$.

⁸ Two earlier papers examine the effect of inheritances on labor supply, but do not focus on retirement. As in our approach, Joulfaian and Wilhelm (1994) use data on inheritance expectations as well as inheritance receipt (from the Panel Study of Income Dynamics), but their data do not identify the recipient within the household. Holtz-Eakin et al. (1993) use a sample of estate tax returns, limited to individuals who received an inheritance. They study labor force participation, and find that receipt of a large inheritance (350 k in 1982 dollars) reduces prime age LFP by 11 percentage points. The generalizability of their results is unclear, because most bequests do not require filing an estate tax return, so their sample is not representative of the population.

⁹ For clarity and simplicity, henceforth we use the term commitment to refer to full commitment ability (enforceable ex ante efficient contracts). Inability to fully commit encompasses limited commitment, where contracts are renegotiated only when a participation constraint binds, and no commitment, with contracts renegotiated every period. Our test cannot distinguish between limited and no commitment, so we lump them together and refer to them jointly as inability to commit, as in Lise and Yamada (2014).

¹⁰ To illustrate this point more transparently, suppose that the inheritance probability distribution takes the following very simple form: with probability π_i , spouse i receives an inheritance of amount I , and with probability $1 - \pi_1 - \pi_2$ neither spouse receives an inheritance ($\pi_1 + \pi_2 < 1$). In this setup, at most one spouse receives an inheritance, and the amount of the inheritance is the same regardless of which spouse receives it. The Pareto weights then can be written as $\mu_i(\pi_i)$, with μ_i increasing in π_i , illustrating the point that a greater likelihood of receiving an inheritance increases bargaining power. Note that inheritance *realizations* are not distribution factors in the commitment model because they are unknown at the time the marriage is formed.

¹¹ This restriction is not imposed in the empirical analysis. If the outside option is to remain married but interact non-cooperatively, we would expect the solution to be similar to this case.

A binding participation constraint causes the period-2 Pareto weight to differ from the period 1 value. If person 1 has a binding participation constraint, then her contribution to the maximand above can be written $M_1 u_1(c_{12}, h_{12}) - \lambda_1 u_{12}^*$, where $M_1 = \mu_1 + \lambda_1$ is the period 2 bargaining power of spouse 1. If spouse 1 receives an inheritance shock large enough to cause her participation constraint to bind, the original “contract” is renegotiated so that her bargaining weight increases by enough (specifically, by λ_1) to make her indifferent between remaining in the marriage and choosing the outside option. In the no-commitment case, receiving an inheritance shock that is large enough to cause a participation constraint to bind causes a shift in resources toward the recipient, resulting in what is effectively a wealth effect, since the inheritance does not alter any relative prices. If leisure is a normal good, we expect this to cause a decrease in hours worked (and participation). Things are more complicated for the non-recipient because there are offsetting effects: (1) his bargaining power declines, so he loses some control over resources, and (2) household wealth increases, so he gains a share of the additional resources available to the household, thanks to resource pooling.¹² The model does not predict which effect dominates. If the latter effect dominates and leisure is highly weighted in person 2's preferences, the decline in his hours of work could be larger than the decline in person 1's hours.

We can write a regression function for period-2 hours of work for spouse i , omitting a household subscript:

$$h_{i2} = \beta_{i1} w_1 + \beta_{i2} w_2 + \alpha_{i1} I_1^* + \alpha_{i2} I_2^* + \gamma_i A_1 + g_i(f(I_1, I_2)) + \varepsilon_{i2}$$

where g_i is a function of ex ante inheritance expectations. It is important to control for inheritance expectations, since they will naturally co-vary with inheritance realizations.¹³ The testable implications of commitment are $\alpha_{i1} = \alpha_{i2}$, $i = 1, 2$: inheritance shocks affect labor supply, but the identity of the recipient of the inheritance does not matter.¹⁴ The inability to commit implies that $\alpha_{i1} \neq \alpha_{i2}$.¹⁵ We expect the “own inheritance” effects to be larger than “spouse inheritance” effects, so we test the null hypothesis of commitment, $\alpha_{i1} = \alpha_{i2}$ against the alternative $|\alpha_{i1}| > |\alpha_{i2}|$. Although this is not a prediction of the theory, in collective household models it is assumed that each individual cares for the spouse but not as much as they care for themselves (Browning et al., 2014, p. 82).

We have assumed egoistic preferences (no externalities in utility), but the result generalizes to any form of non-separable preferences, such as caring preferences and leisure complementarity. The equalities implied by commitment will hold with non-separable preferences,

¹² The inheritance realization is an argument of the Pareto weighting functions, but also enters the period-2 budget constraint. This appears to violate the condition for a variable to be a distribution factor. However, the formulation of the model as described here is equivalent to a formulation in which each spouse has a separate savings account in addition to the couple's joint account, and inheritances are deposited in the individual account of the recipient rather than the joint account. In this formulation, the inheritance realization does not enter the joint budget set, which is the condition for a variable to be a distribution factor (Chiappori et al., 2002). Separate accounts are irrelevant in the commitment case. In order keep the no-commitment analysis comparable to the commitment analysis, we use the joint-account formulation. See Mazzocco (2007) and Voena (2015) for discussion of joint versus individual accounts in household bargaining models with limited commitment.

¹³ In the empirical analysis we estimate this specification as well as a more restrictive version that combines inheritance expectations and realizations into a single unanticipated inheritance variable, $I_i^t = I_i - E(I_i)$.

¹⁴ This result holds in more general models as well. For example, consider a strategic bequest model in which spouse 1 may be able to increase the expected value of her inheritance by providing services such as personal care to the benefactor. If this imposes a cost on spouse 1, for example by reducing her available time to allocate between leisure and employment, then her bargaining power at the beginning of the union would be higher than in the absence of such a consideration. But realization of the inheritance would not alter bargaining power. Similar logic applies if a specific bequeathable good such as a parent's home has sentimental value to one spouse but not to the other. If the recipient plans to keep the parent's home indefinitely after inheriting it, this will reduce the impact of the inheritance on labor supply of both spouses.

¹⁵ It would be interesting to decompose the estimated total effect of an inheritance into the bargaining power and wealth effects, but unfortunately this requires knowledge of initial wealth (A_0), which we do not have.

because under commitment inheritances have only a wealth effect. Thus regardless of the form of preferences the restriction $V_{i2}(\Phi, I_1^*, I_2^*) = V_{i2}^*(\Phi, I_1^* + I_2^*)$ holds under commitment, because this is determined purely by the budget constraint. And this restriction generates the testable implications of commitment. It is worth emphasizing that the empirical analysis does not impose any of the restrictive assumptions used in this section; these assumptions serve to illustrate the main point in as simple a setting as possible.

3. Data

We use data from the Health and Retirement Study (HRS), a national biennial panel study of older individuals and their spouses.¹⁶ The HRS contains an abundance of information on demographic characteristics, health, labor supply, income, and wealth. Our sample includes data from the original HRS cohort born from 1931 to 1941 and interviewed beginning in 1992, the “War Baby” cohort born from 1942 to 1947 and interviewed beginning in 1998, and the Early Baby Boom cohort born from 1948 to 1953 and interviewed beginning in 2004.¹⁷ We use a sample of married couples in which neither spouse was previously married. Limiting the sample to first marriages allows us to focus on the subpopulation for which full commitment is a priori most plausible. Individuals who have experienced a divorce have less stable subsequent marriages on average (Bruze et al., 2015). This may indicate a predisposal toward inability to fully commit. We discuss below how the results differ for couples in which at least one spouse previously divorced. In order to ensure stability of households across survey waves, we keep only couples whose marriage was in progress at the previous survey wave.¹⁸ We examine labor supply behavior in survey years 1996 through 2008.¹⁹ The final analysis sample has 27,448 person-wave (13,724 couple-wave) observations on 3667 married-couple households in which both spouses are between the ages of 45 and 70.²⁰

Our primary measure of labor supply is LFP status at the survey date. We also examine other outcomes such as indicators for currently working for pay and working full-time year-round, and weekly and annual hours worked.²¹ As shown in Appendix Fig. A, LFP rates for the husbands and wives used in our sample follow the typical age profile, with declines beginning in the early 50s and accelerating sharply after age 60. The profiles of husbands and wives are parallel and differ by about 0.2.

The key explanatory variable is receipt of an inheritance since the previous interview (interviews are two years apart on average). The HRS survey asks one member of the household, designated the financial respondent, to answer questions about all inheritances received by the household. If the household received an inheritance from a parent or

¹⁶ Specifically, we make use of the RAND HRS data file (version L), a user-friendly cleaned and processed subset of the HRS data. For certain variables not included in this data file (e.g. inheritance receipts and source of inheritances), we use the RAND enhanced Fat Files datasets. See the RAND and HRS websites for more detail.

¹⁷ We drop a very small number of observations with census region missing or equal to 11 (“Not US/inc US Terr”), and a small number of same-sex couples.

¹⁸ As noted in the previous section, in the absence of commitment ability, receiving a large inheritance could lead to (an efficient) divorce. In this case, the sample of intact marriages would be selected on the basis of an unobserved variable (marital quality) that could be correlated with the degree of intrahousehold cooperation. In practice this is not a concern because <1% of couples divorce between waves. Including these couples in the sample has very little impact on the results.

¹⁹ Data on inheritance expectations are available in the HRS in years 1994 through 2006. We control for lagged inheritance expectations, so we estimate models using data for survey years 1996 through 2008.

²⁰ The HRS Consumption and Activities Mail Survey (CAMS) records information on household expenditure that could in principle be useful for analyzing the effect of inheritance receipt on demand for “assignable” goods such as male and female clothing (see for example Lundberg et al., 1997). Unfortunately, the sample size is too small to be useful.

²¹ Specifically, the respondent is categorized as a labor force participant (LFP = 1) if he or she has full or part time employment, is unemployed, or is partially retired. The respondent is categorized as not in the labor force (LFP = 0) if he or she is retired, disabled, or “not in LF”. Our measure of employment status is based on the variable RxLBRF in the RAND HRS data set, and the alternative measures of hours worked are based on the variables RxJHOURS and RxJWEEKS.

Table 1
Inheritance receipts and expectations.

Variable	N	Mean	p50	p90	p95	p99
<i>(a) Inheritance receipt indicators</i>						
Received over period of observation	3667	0.165				
Received since previous survey	13,724	0.054				
Received since previous survey - husband	13,724	0.023				
Received since previous survey - wife	13,724	0.022				
Received since previous survey - unknown source	13,724	0.010				
Missing: received since previous survey	13,724	0.006				
Missing: amount received	13,724	0.002				
<i>(b) Conditional inheritance amounts (1000s)</i>						
Received since previous survey	714	111.7	41.0	215.0	479.3	1244.7
Received since previous survey - husband	296	123.0	36.8	222.7	530.7	1281.5
Received since previous survey - wife	299	117.4	50.0	282.6	516.9	1281.5
Received since previous survey - unknown source	134	61.4	25.0	138.0	191.1	649.9
<i>(c) Inheritance expectations (lagged 1 period): husband</i>						
Probability of receiving an inheritance >0	13,724	0.353				
Conditional inheritance probability (%)	4839	55.4	50.0	100.0	100.0	100.0
Conditional expected inheritance (1000s)	4405	107.4	35.9	206.8	413.6	1134.1
Missing: probability of receiving	13,724	0.152				
Missing: exp. inheritance amount	13,724	0.030				
<i>(d) Inheritance expectations (lagged 1 period): wife</i>						
Probability of receiving an inheritance >0	13,724	0.415				
Conditional inheritance probability (%)	5700	59.2	60.0	100.0	100.0	100.0
Conditional expected inheritance (1000s)	5067	91.6	35.9	197.7	299.5	823.6
Missing: probability of receiving	13,724	0.049				
Missing: exp. inheritance amount	13,724	0.044				

Notes: The sample has 27,448 person-wave (or 13,724 household-wave) observations on 3667 married-couple households. Dollar amounts are deflated by the Consumer Price Index to the year 2005.

sibling of the financial respondent then we assign the inheritance to the financial respondent. If the financial respondent reports that the household received an inheritance from a parent of his or her spouse, then we assign the inheritance to the financial respondent's spouse. Other responses to the question on the source of an inheritance do not provide enough information to permit the inheritance to be assigned to a particular spouse.²² Specifically, if the inheritance is received from an "other relative", "other individual", "ex-spouse/partner", or the source is missing or unknown, then we do not know whether the inheritance accrued to the husband, wife, or another household member.

Table 1 summarizes the incidence and distribution of inheritance receipt among households in our sample. The first panel shows that 16.5% of couples received at least one inheritance during the 12 year period of observation, and 5.4% received an inheritance between a given pair of interviews, on average. Husbands and wives are about equally likely to be the recipient, and in 19% of cases (0.010/0.054) the recipient cannot be determined.

The second panel in Table 1 shows that the distribution of inheritances amounts is highly skewed. Among the 714 couple-wave observations where at least one inheritance was received, the mean inheritance amount is \$111,700 while the median is \$41,000. Inheritances at the upper tail of the distribution are quite large – the 99th percentile is \$1.24 million.

As discussed earlier, in order to interpret inheritances as shocks, we must control for inheritance expectations. An innovative feature of the HRS is that survey respondents are asked a number of questions about their expectations of future events, including inheritances. The expectations are based on a series of questions asked of *each respondent* (financial and non-financial).²³ Respondents are first asked to rate their chances of receiving an inheritance within the next 10 years, from 0 to 100%. Respondents who report a positive probability are asked how

large the inheritance is expected to be.²⁴ Panels c and d of Table 1 summarize inheritance expectations. Over all person-wave observations in our sample, 35% of husbands and 42% of wives report a positive probability of receiving an inheritance, and conditional on being positive the mean probabilities are 0.55 and 0.59, respectively. Conditional on expecting an inheritance, for both husbands and wives the median expected amount is \$35,900, and the mean expected amount is \$107,400 for husbands and \$91,600 for wives. In 15% of cases the husband does not answer any of the inheritance expectation questions, and in an additional 3% of cases the husband reports a probability of inheritance receipt but does not report an expected inheritance amount; we keep these cases in the sample and include missing-data indicators in the model. We do the same for wives, 5% of whom have missing data on expectations, and an additional 4% of whom report a probability but not an expected amount. Further description of inheritance expectations separately for households that received and did not receive an inheritance is provided in Appendix Table A.

Inheritance expectations are positively correlated with inheritance realizations, as previously reported by Brown et al. (2010). This is illustrated in Appendix Table B. For example, among households that report a zero probability of receiving an inheritance (as of two years prior to the first year the household is included in our sample), only 8% actually received an inheritance during our period of observation. In contrast, among those reporting a 100% probability, 40% actually received an inheritance. Conversely, Appendix Table A shows that 55% of husbands and 64% of wives who received an inheritance reported a non-zero ex ante probability of receipt, compared to 30 and 35% of those who did not receive an inheritance. Further, expected inheritance amounts are correlated with actual amounts received, as illustrated in Appendix Table C.²⁵ Nevertheless, inheritance expectations explain only a small part of the variation

²² Respondents can report up to three inheritances in a given wave. We use the sum of inheritances received from each source.

²³ The inheritance expectation question changed slightly in 2006. Before 2006 the question asks about "... the chances you will receive an inheritance...", while in 2006 the question asks about "... the chances you [or your spouse] will receive an inheritance..."

²⁴ Respondents who do not report a specific value are asked a series of questions that bracket (i.e. assign a lower and upper bound to) the expected inheritance amount. We set the expected inheritance amount equal to the midpoint of the bracket for these respondents.

²⁵ Appendix Table B is very similar to Table 1 from Brown et al. (2010), while Appendix Table C is very similar to Table 4 from Brown et al. (2006).

Table 2
Estimated effects of inheritance receipt, alternative specifications.

	Men			Women		
	Self	Spouse	Other	Self	Spouse	Other
<i>(a) Alternative measures of labor supply</i>						
Labor force participation	−0.040** (0.020)	−0.004 (0.021)	−0.089*** (0.031)	−0.038* (0.022)	−0.005 (0.018)	0.036* (0.022)
Currently working for pay	−0.046** (0.021)	0.006 (0.021)	−0.097*** (0.032)	−0.038* (0.022)	−0.021 (0.020)	0.046** (0.022)
Year round	−0.047* (0.024)	0.007 (0.022)	−0.121*** (0.034)	−0.042* (0.023)	−0.023 (0.022)	0.032 (0.029)
Full time	−0.032 (0.025)	−0.018 (0.023)	−0.073** (0.035)	−0.060** (0.023)	0.006 (0.021)	0.012 (0.035)
YRFT	−0.044* (0.025)	−0.010 (0.023)	−0.095*** (0.035)	−0.060*** (0.023)	−0.006 (0.021)	−0.003 (0.034)
Weekly hours of work	−1.845* (1.013)	−0.300 (0.914)	−3.986*** (1.514)	−2.705*** (0.797)	−0.412 (0.693)	1.368 (1.004)
Annual hours of work	−112.567** (52.148)	−16.756 (46.595)	−222.844*** (77.424)	−112.007*** (39.220)	−25.039 (32.756)	65.030 (50.270)
<i>(b) Alternative samples</i>						
Strong LF attachment	−0.038 (0.025)	−0.007 (0.022)	−0.103*** (0.037)	−0.037 (0.028)	0.005 (0.023)	0.061** (0.026)
Include previously married	−0.024 (0.016)	−0.016 (0.017)	−0.069*** (0.025)	−0.037** (0.018)	−0.006 (0.014)	−0.008 (0.021)

Notes: Ordinary Least Squares (OLS) estimates on the binary indicators of inheritance receipt are reported, with standard errors clustered at the household level. Each row shows results from alternative specifications. The first three columns show results from the regression run on the sample of husbands, and the last three columns show results for the sample of wives. For all of the regressions in panel (a), the sample size is 13,724 person-wave observations. In panel (b), the sample with Strong LF attachment has 9821 observations for men, and 8277 for women. The sample that includes couples with any previously married individuals has 21,514 observations, for both men and women. All regressions include controls for lagged inheritance expectations, lagged inheritances, lagged net worth, pension and Social Security variables, health insurance coverage, health, age, self-employment status, subjective mortality expectations, and whether the respondent's parent died since the previous interview. See Table 3 and Appendix Table E for parameter estimates on the control variables for the specification with Labor Force Participation as the dependent variable. One, two, and three asterisks indicate that the coefficient estimate is significantly different from zero at the 10%, 5%, and 1% level, respectively.

in inheritance receipt, typically around 3%, suggesting either inherent unpredictability of the timing of inheritance receipt, or some form of measurement error in inheritance expectations. However, it is not clear how to think about measurement error in self-reported subjective expectations. There is considerable bunching at focal points in subjective expectation reports, for the probability as well as the amount, but this is likely to reflect uncertainty about probabilities rather than reporting error per se. Hurd et al. (1998) and others provide evidence that subjective expectations reports have considerable informational content despite substantial bunching at focal point responses.

Another key control variable is household wealth at the date of the previous interview.²⁶ We emphasize that inheritances received since the previous interview are *not* included in lagged net worth. Social Security and employer pensions are important sources of wealth that are of particular relevance for older workers. Unlike net worth, claims on Social Security and pensions are illiquid and cannot necessarily be treated as equivalent to other assets. Nevertheless, we follow the conventional approach in the literature and use measures of the expected present discounted value of future Social Security and pension benefits.²⁷ Details of the calculations are described in Appendices A, B, C, and D.

We control for a large number of other variables that may affect LFP and could in principle be correlated with inheritance shocks. These include lagged LFP of the individual and the spouse, lagged self-employment status of both spouses, and whether the employer provides health insurance coverage, both with and without retiree benefits.

²⁶ We use the variable HxATOTA from the RAND HRS dataset, which measures total household net worth. This variable is built up from responses to questions about many types of assets, and incorporates extensive imputations based on partial (bracketed) responses. Note that wealth is measured at the household level in the HRS; the survey does not attempt to identify individual versus joint ownership of each asset.

²⁷ In the scenarios involving continued work beyond the current period, we assume annual earnings are equal to the average of the most recent five years of earnings up to the current period. For these and the additional variables described in this paragraph, we assume standard life table mortality, a 3% real rate of interest, and zero real wage growth. The main specification includes SSW under scenarios (1) and (2), and the increment to SSW from scenario (3) relative to scenario (1).

Other controls include categorical indicators for educational attainment, ethnicity (Hispanic), race, census division, year fixed effects, age (cubic plus dummies for 62–64 and 65+), health status, and recent changes in health status.²⁸ Descriptive statistics are provided in Appendix Table D. It is worth noting that the main results reported below are very robust to alternative specifications of the control variables, including dropping all of the control variables. The only exception is lagged LFP, not surprisingly given that the parameters in a dynamic specification have a different interpretation than in a static specification.

Finally, the death of a parent may have direct effects on labor supply in addition to any effects that operate via receiving an inheritance. For example, if an individual was providing assistance to a parent who then passes away, this could free up time and change the labor supply choice. Alternatively, if the death of one parent leaves the surviving parent in need of assistance, the adult child may need to step in, adding a new constraint on time use. These and other possibilities make it important to control for the death of a parent. The HRS provides information on the vital status of each parent, so we are able to include an indicator for the death of a parent since the previous interview. Any direct effects of parental death on labor supply will be picked up by this variable, alleviating concern that receiving an inheritance might affect labor supply through channels other than the bargaining power and wealth effects we hypothesize.²⁹

²⁸ Theory implies that the wage rates of the individual and spouse should be included in the specification. However, we do not observe a wage rate for non-workers so wage rates are omitted. We estimated several alternative specifications incorporating the wage rate, using a variety of approaches to address the problem of missing wages for non-workers. The estimates from these specifications are virtually identical to those reported below.

²⁹ As a further check we included the number of hours of assistance provided by an individual to his parents, as well as the number of hours provided by the spouse to his or her parents. If time spent caring for an elderly parent is freed up as a result of the death of the parent, we might mistakenly attribute changes in labor supply to receipt of an inheritance, when in reality the change in labor supply was caused by the reduced demand for the individual's time in caring for the parent. An operative strategic bequest motive could also lead to correlation between the amount of care and the size of the subsequent bequest. Controlling for these variables left the results unchanged to the second significant digit in all specifications.

Table 3
Effects of inheritances on labor force participation.

Selected covariates	Men	Women
<i>Inheritance recipient</i>		
Self	−0.040** (0.020)	−0.038* (0.022)
Spouse	−0.004 (0.021)	−0.005 (0.018)
Unknown	−0.089*** (0.031)	0.036* (0.022)
<i>Own inheritance expectations</i>		
Percent chance	0.0002** (0.0001)	−0.00003 (0.00010)
% chance missing	0.005 (0.010)	−0.012 (0.013)
Amount expected	−0.0004 (0.0015)	0.003 (0.002)
Amount missing	−0.019 (0.016)	0.009 (0.015)
<i>Spouse inheritance expectations</i>		
Percent chance	−0.005 (0.007)	0.004 (0.007)
% chance missing	−0.024* (0.014)	−0.010 (0.009)
Amount expected	0.002 (0.002)	−0.000005 (0.000018)
Amount missing	−0.008 (0.014)	−0.006 (0.016)
<i>Labor force status</i>		
Own LFP (one lag)	0.563*** (0.011)	0.564*** (0.011)
Own LFP (two lags)	0.128*** (0.010)	0.129*** (0.010)
Spouse LFP (one lag)	0.038*** (0.009)	0.044*** (0.009)
Spouse LFP (two lags)	−0.005 (0.009)	−0.011 (0.009)
R squared	0.553	0.572
P Value, one-sided test of commitment	0.106	0.111

Notes: Selected coefficient estimates for the sample of men are in the first column, and for the sample of women in the second column. In both cases the sample size is 13,724 household-wave observations. The alternative hypothesis for the one-sided test of commitment is that the coefficient estimate on inheritances received by self is greater than by spouse in absolute value. See Table 2 for additional notes.

In order to test whether this specification is rich enough to adequately control for unobserved factors that could be correlated with inheritance receipt, even after controlling for expectations, we estimated several regression models using this specification to explain outcomes that were determined *before* receipt of an inheritance. If we have adequately controlled for expectations and other factors, inheritance receipt should be uncorrelated with predetermined outcomes. The predetermined outcomes we examined included lagged inheritance receipt, twice-lagged expectations, twice-lagged wealth, and twice-lagged LFP, for the individual and the spouse. We found that this “placebo test” failed for several predetermined variables, indicating that self-reported expectations are not sufficient to eliminate unobserved heterogeneity. As a result, we added the variables described above to the regression specification as additional controls.³⁰

³⁰ The pattern of results for the effect of inheritances on predetermined labor supply did not suggest anticipation effects, where labor supply declines in anticipation of an inheritance before actual receipt. Rather, an inheritance received by a man is associated with higher labor supply by the spouse prior to receipt of the inheritance. The fact that we find selection on observables suggests that there could be selection on unobservables as well. The well-known approach of Altonji et al. (2005) of using the magnitude of selection on observables as a gauge of the potential effects of selection on unobservables does not work in this case because the approach admits only a single potentially endogenous “treatment,” while we have at least two. And, as noted above, the estimates are very robust to alternative sets of controls.

Table 4
Alternative estimates, sample stratified by lagged LFP.

Selected covariates	Men	Women
<i>(a) Alternative sample: Lagged LFP = 1</i>		
Self	−0.035 (0.025)	−0.043 (0.031)
Spouse	−0.019 (0.023)	0.007 (0.026)
P value, one-sided test of commitment	0.319	0.104
<i>(b) Alternative sample: Lagged LFP = 0</i>		
Self	−0.071*** (0.023)	−0.021 (0.024)
Spouse	0.051 (0.043)	−0.026 (0.024)
P value, one-sided test of commitment	0.006	0.562

Notes: Sample sizes for the sample with Lagged LFP = 1 are 9297 for men and 7617 for women, and for Lagged LFP = 0 are 4427 for men and 6107 for women. Estimates for inheritance receipt from unknown source are not shown. See Tables 2 and 3 for additional notes.

4. Results

4.1. Specification

We estimate models of the form

$$y_{it} = \beta_0 + \beta_1 I_{it} + \beta_2 I_{jt} + \beta_3 I_{ut} + \beta_4 E_{it-1} + \beta_5 E_{jt-1} + \beta_6 X_{it-1} + \beta_7 X_{jt-1} + \beta_8 Z_i + \varepsilon_{it}$$

y_{it} is a binary indicator of LFP of member i in a couple (the couple subscript is omitted). I_{it} is an indicator of inheritance receipt by individual i since the previous interview, I_{jt} is an indicator for inheritance receipt by i 's spouse since the previous interview, and I_{ut} is the inheritance indicator for cases in which the recipient within the household cannot be determined (unknown). We focus on whether an inheritance was received, rather than the amount, because the amount received is highly skewed and is probably measured with substantial error. We report results for specifications using amounts inherited below.

E_{it-1} and E_{jt-1} are the inheritance expectations of the individual and spouse as of the previous interview. X_{it-1} and X_{jt-1} are vectors of time-varying spouse-specific variables, including two lags each of own and spouse LFP, inheritance receipt, inheritance expectations, and household wealth, and Z is a vector of fixed household characteristics. Conditional on inheritance expectations as of the previous interview (and the other control variables), we interpret β_1 , and β_2 as the effects of inheritance shocks.

The main estimates are from linear probability models.³¹ Estimates are presented separately for husbands and wives. We include all couples regardless of labor force participation in the previous wave. This makes it possible to capture effects of inheritances on reentry to the labor force as well as exit. It is well known that retirement patterns can be complex, with repeated exit and entry, so looking only at exit could miss part of the impact of an inheritance shock. We also present results that split the sample according to lagged LFP.

4.2. Main results

Table 2 presents selected results from OLS estimates of the model for alternative labor market outcomes and samples. The first three columns shows results for men, with each row in Panel a reporting results for a different labor market outcome. The last three columns show results for women. Coefficient estimates on the variables other than those shown in the table are reported for the specifications in the first row in Appendix Table E. The standard errors are clustered by household. The first column shows that receipt of an inheritance by a married man causes a 4.0 percentage point (pp) decline in LFP, a 7% effect relative to the

³¹ Marginal effects from probits are very similar to the linear regression estimates.

Table 5
Estimates allowing nonlinear inheritance effects.

Sample	Men		Women	
Recipient	Self	Spouse	Self	Spouse
<i>(a) Continuous inheritance amounts</i>				
Inheritance amount	0.010 (0.007)	−0.015 (0.009)	−0.014 (0.009)	−0.002 (0.007)
<i>(b) Indicators for inheritance amount below or above median</i>				
Above median	−0.047 (0.029)	−0.035 (0.029)	−0.037 (0.028)	−0.022 (0.028)
Below median	−0.039 (0.027)	0.032 (0.029)	−0.054* (0.032)	0.008 (0.022)
<i>(c) Indicators for quartile of inheritance amount distribution</i>				
Quartile 1 (lowest)	−0.096** (0.042)	0.070 (0.043)	−0.067 (0.047)	−0.024 (0.036)
Quartile 2	0.021 (0.030)	−0.005 (0.038)	−0.039 (0.040)	0.044** (0.020)
Quartile 3	−0.041 (0.036)	−0.020 (0.042)	−0.002 (0.037)	0.003 (0.033)
Quartile 4 (highest)	−0.053 (0.043)	−0.050 (0.039)	−0.074* (0.040)	−0.047 (0.046)

Notes: The first two columns report results from three different regressions for men, and the last two columns report results for three different regressions for women. In panel (a), inheritance amounts are measured in units of \$100,000 (deflated to 2005 dollars). In panels (b) and (c), the omitted category is non-recipients. See Tables 2 and 3 for additional notes.

sample mean LFP rate of 0.61 for men. Receipt of an inheritance by the wife causes a small decrease of 0.4 pps in his LFP. The effects of own and spouse inheritance receipt for women shown in the second set of columns are very similar in sign and magnitude to the effects for men.

The effects of inheritances for which the recipient within the household is unknown are surprising; the estimate is quite large in magnitude (8.9 pps) for men, and it is positive for women. A possible explanation is that these “unknown” inheritances may be disproportionately controlled by the husband. The financial respondent is male in roughly 60% of observations, and Laitner and Sonnega (2010) note that inheritances received by non-financial respondents appear to be underreported in the HRS.³²

The next four rows of Table 2 show results for alternative binary measures of LFP, for which the dependent variable is equal to one for currently working for pay, working year-round, working full time, and working year-round full time, respectively. The effects of own inheritances are robust to these alternative measures of employment, with a tendency toward larger effects than in models of LFP. Spouse inheritance effects are generally small, although somewhat larger on average than in the first row. The next set of results is for weekly and annual hours of work. The own inheritance effects are negative, and spouse effects are much smaller in absolute value than the own effects.

The last two rows show parameter estimates from alternative samples. First we limit the sample to men and women with a relatively strong attachment to the labor force: work experience of at least 10 years, job tenure of at least two years (if employed), and out of the labor force for less than five years (if not employed). The inheritance effects are very similar in magnitude to those reported in the first row, and a bit less precisely estimated. Second, we broaden the sample to include couples with previously married individuals. Results are qualitatively similar to those reported in Table 2, although the magnitudes of the own inheritance effects are slightly smaller for men.

³² In the robustness analysis discussed below, we explore alternative approaches to dealing with inheritances that cannot be assigned to a specific household member. It turns out that the results are generally insensitive to alternative approaches, including dropping these cases and assigning them to the husband.

Given the similarity of the results across alternative labor supply measures, we focus on LFP in the remainder of the analysis. Table 3 repeats the estimates for LFP from the first row of Table 2 along with additional estimates from these regressions. Inheritance expectations turn out to have little association with LFP, despite the fact that they are positively correlated with inheritance realizations, as noted earlier. The null hypothesis that the coefficients on the four own inheritance expectations variables shown in Table 3 are jointly equal to zero cannot be rejected at the 10% level of significance. There are a few statistically significant individual coefficient estimates, but the implied magnitudes are very small.

One way to illustrate the magnitude of the inheritance receipt effects is to compare them to the estimated effects of other events. For example, Table E in Appendix A shows that the impact on LFP of a self-reported deterioration in health since the previous interview is −0.038 for men and −0.036 for women. Thus the own-inheritance effects are comparable in magnitude to the effect of deterioration in health. This suggests a relatively large impact of inheritance receipt in view of the importance of health declines for LFP.

Our estimates are roughly in line with earlier studies examining the effects of inheritances on labor supply, despite differences in sample composition and time period. Brown et al. (2010) report household-level inheritance effects of about 0.02 on the probability of exit from the labor force (Table 2), based on a sample of older individuals and married couples in the HRS.³³ The estimates of Holtz-Eakin et al. (1993) imply that a \$100,000 inheritance would have caused LFP to fall by 0.039 among prime age men and women during the 1980s. Joulfaian and Wilhelm also analyze labor supply of prime age men and women during the 1980s, and find smaller effects: a \$100,000 inheritance would have caused a reduction in annual hours of work of 24 for women, and less for men. We can also compare our estimates of the inheritance effects to estimates of the effects of lottery winnings on labor supply, an arguably similar type of shock to wealth. Using a measure of labor force participation defined by an annual earnings threshold of about \$4000, Cesarini et al. (2015) estimate that winning a \$100,000 lottery prize would cause LFP to fall by 0.013 among a sample of Swedish men and women ages 21 to 64. Imbens et al. (2001) estimate a marginal propensity to earn from lottery winnings of −0.167 for 55–65 year old winners in the 1980s. If the change in earnings were due entirely to changes in LFP, holding hours worked constant, winning a \$100,000 lottery would cause LFP to decline by 0.052.³⁴ Overall, our results are well within the range of estimates from the previous literature.

The *p*-value for the one sided test of commitment is shown at the bottom of the columns in Table 3. The test is for the equality of the

³³ Brown et al. (2010) also report estimates using a long-difference sample, with one observation per household summarizing labor force exits between 1994 and 2002. These estimates are not directly comparable to ours.

³⁴ Imbens et al. estimate a model of annual earnings of the form $y = \beta_0 + \beta_1 L/20$, where L is the total payout, which is spread over 20 years. Their estimate of β_1 is −0.0167 for ages 55–65. Let $E(y) = E(wH)$, where w is the hourly wage, H is annual hours worked, E is the expectations operator, and L is a dummy for LFP. If the lottery only affects L , then $dy/dL = \beta_1/20 = wHdE(L)/dL$, so $dE(L)/dL = \beta_1/(20wH) = -0.00000521$, given the (overall sample) mean value of wH of \$16,000. Multiply this by \$100,000 to get the effect of winning a \$100,000 lottery on LFP: −0.0521. Cesarini et al. estimate a regression of the form $L = \beta_0 + \beta_1 L$, with L measured in units of one million Swedish Krona and L measured by whether annual earnings exceed \$25,000 Krona. The coefficient estimate is −2.067 (Table 4). 1,000,000 Krona is equivalent to \$160,000, so the effect of winning a \$100,000 lottery is $(-2.067)/(0.16 \times 100) = -0.013$. Holtz-Eakin et al. estimate a logit model of transitions from employment to non-employment for a sample of unmarried inheritance recipients, using a quadratic in the inheritance amount. From their Table 3 (column 3), we computed the marginal effect of a \$100,000 inheritance as −0.039, evaluated at $L = \$164,000$ and mean LFP = 0.82. The median inheritance in their sample was \$82,000 in 1982 dollars (we computed this as a weighted average of medians reported by group in their Table 1). We doubled this to account for inflation since 1982. Joulfaian and Wilhelm computed the estimates cited in the text based on their Table 5 for women and Table 3 for men.

Table 6

Estimates of the effect of inheritance receipt relative to expectations.

	Men	Women
<i>(a) Inheritance relative to expectations</i>		
Self	−0.011 (0.015)	−0.012 (0.015)
Spouse	0.017 (0.011)	0.015 (0.011)
P value, one-sided test of commitment	0.062	0.076
<i>(b) Inheritance relative to expectations, by sign</i>		
Self - positive	−0.034 (0.023)	−0.035 (0.027)
Spouse - positive	0.064** (0.030)	0.030 (0.022)
Self - negative	0.014 (0.010)	0.011 (0.010)
Spouse - negative	0.052* (0.029)	0.031 (0.025)
P Value, one-sided test of commitment		
Positive	0.005	0.033
Negative	0.102	0.224

Notes: The first two columns report results from two different regressions for men, and the last two columns report results from two different regressions for women. See Tables 2 and 3 for additional notes, and see text for detail on measures of inheritances relative to expectations.

coefficients on own and spouse inheritance against the alternative that the own effect exceeds the spouse effect in absolute value. We show results for one tailed tests, although as discussed above the theory does not predict the sign of the difference in coefficients if full commitment does not hold. Nevertheless, the most plausible alternative to equality is that the own effect exceeds the spouse effect in absolute value (which is the case for both husbands and wives as reported in Table 3, and in all results from alternative specifications reported in Table 2). The *p* values are 0.106 and 0.111 for husbands and wives, respectively. For the nine specifications displayed in Table 2, the null hypothesis is rejected at the 10% level in three cases for men and five cases for women (results not shown). As noted above, while inheritances are not rare in this population, they are not very common at a two-year frequency (5%), so lack of sufficient statistical power could be one reason for the less-than-decisive results. Below, we explore ways to increase the precision of the estimates.

4.3. Additional results

(i) Table 4 presents results from estimates that condition on LFP = 1 at the previous interview in the upper panel, and LFP = 0 in the lower panel. For men, the effect of an inheritance on LFP for men initially out of the labor force is twice as large in absolute value as the effect for men initially in the labor force (−0.071 and −0.035), while the effects for women follow the opposite pattern. The null hypothesis of commitment is rejected only for men with lagged LFP = 0.

(ii) We next report results from specifications that use the amount inherited. The first row of Table 5 indicates that the amount inherited (in units of \$100,000) has small estimated effects on LFP of both husbands and wives, regardless of the identity of the recipient. Next, we consider the possibility that effects of the amount inherited may be nonlinear. The remaining rows of Table 5 report results from a specification that includes indicators for whether the inheritance amount is above or below the median (conditional on receiving an inheritance; not receiving is the omitted category), and from a specification with indicators for the quartile of the positive part of the inheritance distribution. The effects of own-inheritance receipt are much larger in these specifications. This suggests that the effect is nonlinear, perhaps explaining the poor results in the linear specification in the first row.

Table 7

Effects of inheritance receipt on LFP, alternative specifications.

	Men	Women
<i>(a) Subjective commitment indicator</i>		
Self × “Agreement”	−0.109 (0.079)	−0.046 (0.044)
Spouse × “Agreement”	−0.197** (0.089)	−0.009 (0.035)
Self × “Disagreement”	−0.047** (0.024)	−0.065*** (0.024)
Spouse × “Disagreement”	0.009 (0.024)	−0.023 (0.021)
P value, one-sided test of commitment		
Agreement	0.777	0.260
Disagreement	0.048	0.091
<i>(b) Fixed effects estimator</i>		
Self	−0.028 (0.020)	−0.041* (0.023)
Spouse	−0.018 (0.022)	0.011 (0.021)
P value, one-sided test of commitment	0.374	0.039
<i>(c) Cumulative inheritance effects, three periods (six years)</i>		
Self	−0.135*** (0.049)	−0.093*** (0.056)
Spouse	−0.010 (0.053)	−0.036* (0.049)
P value, one-sided test of commitment	0.038	0.213

Notes: The estimates in panel (a) are from a specification where the binary indicator of inheritance receipt is interacted with a subjective commitment indicator set to “agreement” if the husband and wife both agree on who has the most say in major decisions, both say that the time they spend with their spouse is “extremely enjoyable”, and both say that they like to spend their free time “together”. Otherwise the indicator is set to “disagreement”. The estimates in panel (b) are from a specification that includes individual fixed effects. Panel (c) presents the estimated cumulative effect of inheritances on labor force participation, based on the estimates in Table 3 and Appendix Table E. See Tables 2 and 3 for additional notes.

This specification also reduces the impact of measurement error. There is no evidence of a dose response, so the binary indicator specification seems appropriate.³⁵

(iii) The main specification used in this paper does not impose a tight relationship between the effect of expectations and actual inheritance receipt. An alternative approach is to construct a measure of unanticipated inheritances, $I_{it} - E_{it-1}$, as the main explanatory variable. This imposes more structure but is perhaps easier to interpret. The regression specification in this case is

$$y_{it} = \beta_0 + \beta_1(I_{it} - E_{it-1}) + \beta_2(I_{jt} - E_{jt-1}) + \beta_3I_{ut} + \beta_3X_{it-1} + \beta_4X_{jt-1} + \beta_5Z_i + \varepsilon_{it}.$$

In order to construct such a measure we have to convert expectations, which are measured over a ten-year horizon, to the same horizon as actual receipt, which is measured since the previous interview, two years on average. If we assume that the respondent's subjective probability of receiving an inheritance is the same in each of the five two year periods covered by her response to the ten year expectation question, then the probability of receiving an inheritance in the next two years is $q = 1 - (1 - p)^{1/5}$, where p is the ten year probability. We use $I_{it} - q_{it-1}$ as the “surprise.”

³⁵ Another approach to investigating the impact of measurement error is to trim outliers from the sample, to determine whether exceptionally large inheritance values might be erroneous. Table F in Appendix A shows that as successive top percentiles are dropped, the effect of the own amount received on LFP increases monotonically down to the 96th percentile. This suggests that measurement error in the form of misreported large values of the amount inherited is important.

Table 8
Estimated effects, alternative treatment of inheritances from unknown source.

Selected covariates	Men	Women
<i>(a) Unknown enters separately</i>		
Self	−0.040** (0.020)	−0.038* (0.022)
Spouse	−0.004 (0.021)	−0.005 (0.018)
Unknown	−0.089*** (0.031)	0.036* (0.022)
<i>P</i> value, one-sided test of commitment	0.106	0.111
<i>(b) Combine unknown with male</i>		
Self	−0.059*** (0.017)	−0.038* (0.022)
Spouse	−0.004 (0.021)	0.004 (0.014)
<i>P</i> value, one-sided test of commitment	0.021	0.045
<i>(c) Drop unknown</i>		
Self	−0.042** (0.020)	−0.039* (0.022)
Spouse	−0.005 (0.021)	−0.011 (0.018)
<i>P</i> value, one-sided test of commitment	0.104	0.155
<i>(d) Drop if inheritance not received from parent of financial respondent</i>		
Self	−0.052** (0.026)	0.008 (0.033)
Spouse	0.014 (0.035)	−0.028 (0.024)
<i>P</i> value, one-sided test of commitment	0.067	0.822
<i>(e) Combine unknown with non-financial respondent</i>		
Self	−0.058*** (0.020)	−0.021 (0.018)
Spouse	−0.017 (0.018)	0.004 (0.016)
<i>P</i> value, one-sided test of commitment	0.063	0.150

Notes: Sample sizes in panels c and d are 13,568 and 13,280 household-wave observations, respectively. See Tables 2 and 3 for additional notes.

A useful feature of this approach is that it allows us to examine the possibility of asymmetry in response to positive and negative shocks. It seems intuitive that the negative “surprise” of not receiving an inheritance, given a positive subjective probability, may evoke a smaller change in labor supply than would the positive surprise of receiving one, conditional on a subjective probability less than one.

Table 6 reports results from a specification in which the inheritance variables appear in the form of actual receipt [0, 1] minus subjective probability of receipt. The upper panel shows results that do not distinguish positive and negative shocks, as a baseline. The effects of own inheritance receipt relative to expectations are small, but the hypothesis of commitment is rejected at the 10% level of significance. In panel b we expect a negative effect on LFP of a positive surprise and a positive effect of a negative surprise. This pattern is confirmed, and the effect on LFP of a positive inheritance surprise is larger in absolute value than the effect of a negative inheritance surprise. The null hypothesis of commitment is rejected for positive surprises at the 1% and 5% level of significance for husbands and wives, respectively, and is not rejected in either of the specifications for negative surprises.³⁶

³⁶ The results in Table 6 use the same sample as in the previous tables, which include cases with missing inheritance expectations. In previous tables, missing values were replaced by zeros, and dummies for such cases were included. In the actual-minus-expected specification, this may not be an innocuous approach. We re-estimated the models in Table 6 using smaller samples that dropped cases with missing expectations. The results were very similar to those reported in Table 6. These results are available on request.

(iv) An interesting question is whether there are observable subgroups of couples for which commitment is a priori more plausible. We examine this issue using self-reported information about the marriage that may help distinguish between couples more and less likely to be able to agree to self-enforcing contracts. Specifically, we identify three relevant measures in the HRS data: whether the spouses agree on who has the most say in major decisions, whether they both say that the time they spend with their spouse is “extremely enjoyable”, and whether they both say that they like to spend their free time “together”.³⁷ It seems plausible that couples who agree on these statements are more likely to behave as if they have signed a binding contract at the beginning of marriage. However, these measures were collected many years after the beginning of the marriage for most couples, so they may not convey much information about commitment at the time of the marriage.

The results in panel a of Table 7 show some support for differences in commitment according to whether the couple agrees in their responses to the three questions. Among men in couples who disagree with any of the three statements, the null hypothesis of commitment can be rejected at the 5% level of significance for husbands and at the 10% level for wives, while the null cannot be rejected in couples who agree with all the statements.

(v) One issue in interpreting our results is whether the OLS estimates are likely to be biased. We can be reasonably confident in interpreting inheritance receipt as unexpected, given controls for expected inheritances, so bias resulting from anticipation effects is unlikely.³⁸ However, Brown et al. (2010) point out that unobserved heterogeneity is another potential source of bias in estimating the effect of inheritance receipt on labor supply.³⁹ We use a fixed effects estimator to address this issue. This approach is feasible because our main specification pools labor force participants and non-participants, so the effects are identified by both exit and entry. About one third of the sample ever changes labor force status during the period of observation.

Panel b of Table 7 reports results from the fixed effects specification. The estimate for the husband is smaller and less precise than the OLS estimate in Table 3, but the magnitude is fairly similar. The FE estimates for wives are very similar to the OLS estimates, and the null hypothesis of commitment is rejected at the 5% level of significance. Overall, these estimates suggest that any bias in the OLS estimates is likely to be small.

(vi) The results discussed so far are short run effects that occur within two years of receiving an inheritance. Our model is dynamic, so we can trace out longer run impacts as well, accounting for effects transmitted via once-lagged and twice-lagged own and spouse LFP, as well as effects arising from once-lagged and twice-lagged inheritance receipt (see Appendix Table E for coefficient estimates on these variables). The cumulative effects operate through both the husband and wife equations as a consequence of the cross-spouse effects. Panel c of Table 7 shows cumulative effects of inheritance shocks through three periods (six years) based on the estimates in Table 3 and Appendix Table E. The own-inheritance effects are two to three times larger after three periods compared to the one-period effects reported in Table 3, and spouse effects are larger as well. For husbands, the own

³⁷ These measures were previously used by Friedberg and Webb (2006) and Maestas (2001) in analysis of retirement decisions.

³⁸ In this context an anticipation effect could arise if an individual's parent is in very poor health, leading to the expectation of an imminent inheritance. If the individual leaves the labor force before receiving the inheritance, this would be an anticipation effect.

³⁹ For example, more risk-averse individuals are likely to accumulate more wealth than their less risk-averse counterparts, perhaps in part by retiring later. They are more likely to die with positive assets, bequeathed to their children. If characteristics such as risk aversion are correlated across generations within families, then the negative association between inheritance and labor supply could be spurious. Controlling for inheritance expectations may deal with this problem, but we cannot be certain of this.

effect increases by more than the spouse effect, leading to rejection in the test for commitment. However, for women, the spouse effect increases by more than does the own effect, so the commitment hypothesis cannot be rejected in the longer run.

4.4. Robustness to alternative treatments of unknown-recipient inheritances

In the remainder of this section we return to the issue of how to deal with inheritances for which the recipient within the household cannot be identified. Table 8 shows results from five different approaches to dealing with such cases. Panel a repeats results from Table 3, where inheritances from an unknown source enter separately in the specification. As noted above, receipt of an inheritance whose owner within the household is unknown has a large effect on LFP of men, more than twice as large as the own-inheritance effect. This suggests that men may have more control of these inheritances. The financial respondent is male in roughly 60% of observations, which could lead to better reporting of inheritances from relatives of the male. This suggests a specification in which inheritances with an unknown recipient are pooled with inheritances received by men.⁴⁰ Results for this specification are presented in Panel B of Table 8. The own-inheritance effect increases from -0.04 to -0.059 for husbands and is unchanged for wives. In both cases, the null hypothesis of commitment is rejected at the 5% level of significance.

Another approach to dealing with inheritances for which the recipient is unknown is to drop those cases. Panel c reports results using this approach. The results are very similar to those in the first panel, with a slight loss of precision.

Panel d reports results in which only observations with an inheritance received by the financial respondent are used (along with non-recipients, as usual). This is based on the point made by Laitner and Sonnega (2010) that there appears to be significant underreporting of inheritances for non-financial respondents. The results for men are quite similar to those in Panel a, but the coefficient estimates for women change sign and are small in absolute value. This is not surprising given that the majority of financial respondents are men. Many inheritances received by women are omitted in this approach. Laitner and Sonnega (2010) also provide evidence from the HRS that inheritance reporting is generally more accurate when the inheritance is received by the financial respondent, regardless of the recipient's gender. It follows that "unknown" sources of inheritance may be more likely to belong to the non-financial respondent. We use this approach in Panel e, and again, estimates are similar to the preceding panels. Overall, the results are not very sensitive to the treatment of unknown inheritance recipients, and in several cases we reject commitment.⁴¹

5. Conclusions

This paper provides new evidence on the ability of spouses to commit to an ex ante efficient allocation of resources within the household.

⁴⁰ In practice, this means that for the husbands in the sample, inheritances with an unknown recipient are pooled with inheritances received by "self". And for the wives in our sample, inheritances with an unknown recipient are pooled with inheritances received by "spouse".

⁴¹ Gouskova (2013) provides evidence of another source of measurement error: "forward telescoping" in inheritance reporting, a type of memory error when events are remembered as occurring more recently than they actually did. As a result, respondents might report the same inheritance more than once if an inheritance that was received before the previous interview is erroneously recalled as having occurred after the interview. When we estimate our models using corrected measures of inheritances as suggested by Gouskova (2013), the results (not shown) are nearly identical to the main results presented in Table 3.

The analysis exploits unique data from the HRS on inheritance expectations and inheritances received by each spouse in married couples. Controlling for inheritance expectations, we interpret the impact of inheritance receipt as a shock to wealth. The inheritance is legally under the control of the recipient and is not subject to marital property law, so it can be interpreted as a distribution factor: a variable that affects bargaining power within the household but has no direct effect on preferences or the joint budget constraint. The results are inconsistent with the hypothesis that households are able to fully commit to an ex ante efficient contract at the beginning of marriage. Commitment ability has been tested and rejected in other contexts, as discussed in the introduction, so our evidence is consistent with previous evidence on this issue.

The finding that commitment is infeasible, at least for older households, has implications for US Social Security policy, specifically for proposed reforms to spouse and survivor benefits. Under current policy a spouse (typically the wife) who has worked only intermittently may be eligible for a spouse benefit equal to 50% of her husband's benefit, if the spouse benefit exceeds the benefit to which she is entitled based on her own earnings history. Upon the death of her spouse, a woman whose retired-worker benefit is less than the benefit of her deceased spouse is eligible for a survivor benefit equal to 100% of his benefit. Martin (2012) argues that this policy results in many inequities and unintended consequences, and is inconsistent with the community property approach to marital assets now used by all states. Notably, employer-provided pensions are treated as community property, so it is ironic that pensions provided via Social Security are not. One proposal to reform the system would be to combine the earnings received by spouses during marriage and divide them equally between the spouses for purposes of determining Social Security benefits. Because spouses are unable to commit to a long term contract, one consequence of such a reform would be to increase the decision power of low-earning wives.⁴²

A natural extension of the analysis is to exploit inheritance shocks for structural estimation and identification of an intertemporal collective model of the household. Voena (2015) solves and estimates a model without commitment, using divorce and marital property law changes as distribution factors for identification. The advantages of using inheritances as distribution factors are that there is household-level variation rather than only cross-state and time variation for divorce and marital property law, and data on inheritance expectations can be used in a natural way in estimation (see van der Klaauw and Wolpin, 2008, for an example of structural estimation using expectations data).

Acknowledgements

Blau thanks the National Institute on Aging grant R01-AG02199 for support. The data used in this paper are from the Health and Retirement Study, which is sponsored by the National Institute on Aging (grant number NIA U01AG009740) and is conducted by the University of Michigan. We appreciate helpful comments in seminars at OSU, Middle Tennessee State, the Michigan Retirement Research Center, National University of Singapore, Singapore Management University, the 2015 World Meeting of the Society of Labor Economics, the 2015 World Congress of the Econometric Society, and especially from Maria Casanova, Bruce Weinberg, John Ham, Ken Yamada, Eric French, and Robert Moffitt. We are responsible for all opinions and errors.

⁴² Mazzocco (2007) argues that the inability of households to commit to a long term contract makes it possible for public programs to redistribute decision power within the household. Persson (2014) showed that a reform of the survivor's insurance system in Sweden that drastically reduced the dependence of the survivor's benefit on the deceased spouse's earnings history had a substantial impact on divorce and marital selection.

Appendix A

Social Security Wealth

We compute the expected present discounted value (EPDV) of Social Security benefits (Social Security Wealth, or SSW) under several alternative assumptions about labor force exit and claiming: (1) exit in the current period and never return to work, and claim the benefit at 62, (2) exit in the current period and successfully apply for Social Security Disability benefits, (3) work until age 62, exit the labor force at 62, and claim the Old Age and Social Insurance (OASI) benefit at 62 and never return to work, (4) same as (3) but at 65, (5) same as (3) but at 70. These are used in alternative specifications to determine whether the results are sensitive to the specific assumptions. Earnings data from Social Security Earnings Records are used to compute benefits under a set of assumptions about future earnings. We also use these data to construct a measure of the EPDV of remaining lifetime earnings, included as a control variable.

We use administrative Social Security data on annual covered earnings of HRS respondents from 1951 to 2003 for the original HRS cohort, and from 1951 to 1999 for the War Baby and Early Boomer cohorts. In addition, we use annual earnings for the previous calendar year reported by HRS respondents in even-numbered survey years. These earnings, capped at the maximum taxable earnings, are used to extend the earnings series through 2007 and to fill in missing odd years from 1991 to 2003. We assign the average of earnings in adjacent years to fill in earnings for years with missing data. If the last observation on earnings prior to the assumed date of labor force exit in a given scenario (see below for the scenarios) is missing, we assign the value from the preceding year. Pre-1951 earnings were set equal to 1951 earnings. Labor force entry was assumed to occur at age $a = \max\{e + 6, 16\}$, where e is years of schooling completed.

The earnings data are used as input to the Social Security Administration's *anyplab* program to calculate Social Security retirement and disability benefits under a variety of scenarios.⁴³ For each survey wave observation on a given individual (even numbered years from 1996 to 2008) the scenarios include the following: (1) Stop working in the current year, never return to work, and claim at the first date of eligibility (the current year, if already at least 62). (2) Work until age 62, claim at age 62. (3) Work until age 65, claim at 65. (4) Work until 70, claim at 70. (5) Stop working in the current year, apply for SSDI benefits, and begin receiving SSDI benefits in the following year. For scenarios (2)–(4) we do not use observed earnings for the years between the survey year of interest and the assumed year of claiming, because these reflect actual work behavior, and our scenarios assume constant earnings until retirement. We impute earnings for the years between the survey year of interest and the assumed year of claiming using the average of the five most recent years of earnings observed prior to the survey year.

The annual benefits calculated by the *anyplab* program under each scenario are used to compute Social Security Wealth, the expected present discounted value of remaining lifetime benefits, discounted at a 3% rate back to the survey year, using life table mortality schedules to discount for survival risk. We assume that benefits remain constant in real terms after claiming. We use benefits for both the husband and wife to calculate SSW, because benefits are determined at the household level. For each study year observation, we compute the spouse benefit corresponding to the individual's retired worker benefit, and assign to the spouse

the larger of the spouse benefit and her (or his, if the wife's retired worker benefit is larger) own retired worker benefit. We make the following assumptions about when the spouse claims his or her benefit in a given scenario:

- If the spouse's benefit is based on her own earnings (because her retired worker benefit exceeds her spouse benefit) the spouse claims at the same age as the respondent in the scenario: 62, 65, or 70 (which will be in a different year unless the spouses are the same age).
- If the spouse's benefit is based on her husband's earnings record, then she claims at 62 if she is young than her spouse, and she claims when he turns 62 if she is older than her spouse.

We account for joint survival probabilities (both survive to the next year, the husband survives and the wife dies, the husband dies and the wife survives, and both die), and assign a survivor benefit for the cases in which one spouse is assumed to die and the other is assumed to live.

Appendix B

Defined benefit pension wealth

We use employer-reported data on the Defined Benefit pension plans held by HRS respondents. The HRS provides these data along with software that can be used to calculate benefits under alternative scenarios regarding earnings, inflation, interest rates, and the date of claiming. We use a single self-reported annual earnings observation in this case, along with an assumed growth rate of zero, instead of the earnings series described above, for two reasons. First, Social Security earnings are capped at the maximum taxable level, and the cap is binding in many cases. Second, the self-reported earnings variable is built into the calculator, and it is very cumbersome to use the pension calculator with a user-provided earnings series. The pension calculator is used to compute the present discounted value of the annuity to which an individual would be entitled, using an assumed real interest rate of 3% and life table mortality schedules, for each of the first four employment-claiming scenarios described above. A two thirds joint and survivor annuity is assumed, and mortality risk of the spouse is incorporated in the present value calculations.

Appendix C

Defined contribution balances

We use respondent-reported DC balances because relatively few DC plans held by respondents are included in the pension data base. Balances are summed for all plans held by a respondent at a given survey wave. We also used DC balances computed by the pension calculator, as an alternative to respondent-reported balances, with similar results.

Appendix D

Lifetime earnings

We use average capped earnings in the five years prior to the survey year as our measure of average lifetime earnings. As noted above, we assume constant real earnings from a given survey year until the assumed year of retirement, and we use average earnings in the five years prior to the survey year to project earnings forward. We could use observed capped earnings up to the survey year, but this will be equal to zero if the individual actually did not work in one of those years.

⁴³ We use the batch version of the calculator, *anyplab*, available at <http://www.ssa.gov/oact/anyplab/anyplab.html>.

SSW, DB Pension Wealth, and DC Balances are deflated to 2005 dollars using the CPI-U.

Appendix Table A

Inheritance expectations, by sex and whether household ever received an inheritance.

Variable	N	Mean	p50	p90	p95	p99
<i>(a) Husband, household received at least one inheritance</i>						
Probability of receiving an inheritance >0	2961	0.546				
Conditional inheritance probability (%)	1617	64.5	75.0	100.0	100.0	100.0
Conditional expected inheritance (1000s)	1502	123.3	43.4	249.0	497.9	1085.6
Missing: probability of receiving	2961	0.120				
Missing: exp. inheritance amount	2961	0.037				
<i>(b) Husband, household did not receive any inheritances</i>						
Probability of receiving an inheritance >0	10,763	0.299				
Conditional inheritance probability (%)	3222	50.9	50.0	100.0	100.0	100.0
Conditional expected inheritance (1000s)	2903	99.1	30.0	197.7	373.4	1134.1
Missing: probability of receiving	10,763	0.161				
Missing: exp. inheritance amount	10,763	0.028				
<i>(c) Wife, household received at least one inheritance</i>						
Probability of receiving an inheritance >0	2961	0.641				
Conditional inheritance probability (%)	1898	67.6	80.0	100.0	100.0	100.0
Conditional expected inheritance (1000s)	1706	111.7	39.5	217.1	395.3	1033.9
Missing: probability of receiving	2961	0.037				
Missing: exp. inheritance amount	2961	0.061				
<i>(d) Wife, household did not receive any inheritances</i>						
Probability of receiving an inheritance >0	10,763	0.353				
Conditional inheritance probability (%)	3802	55.0	50.0	100.0	100.0	100.0
Conditional expected inheritance (1000s)	3361	81.3	31.1	186.7	263.6	823.6
Missing: probability of receiving	10,763	0.052				
Missing: exp. inheritance amount	10,763	0.039				

Notes: In panels (a) and (c), the samples include only households that receive at least one inheritance over the period of observation, and in panels (b) and (d) the samples include only households that receive zero inheritances over the period of observation. See Table 1 for additional notes.

Appendix Table B

Expected vs realized inheritances.

Probability of receiving inheritance	% of sample	Median conditional value of expected inheritance (\$)	Share that received an inheritance	Median conditional value of inheritances received (\$)
0.00	0.55	–	0.08	35,677
0.01 to 0.49	0.13	13,178	0.18	29,063
0.50	0.09	39,534	0.25	48,063
0.51 to 0.99	0.11	51,694	0.34	52,834
1.00	0.09	65,891	0.40	63,658
Unknown	0.04	–	0.09	29,311
All	1.00	39,534	0.17	45,530

Notes: Realized inheritance receipts are cumulative from the first year the household is included in the sample through 2008. For the HRS, WB, and EBB cohorts the first observation is in years 1996, 2000, and 2006, respectively. The probability of receiving an inheritance is the expectation as of the wave prior to the first observation in our sample (i.e. 1994, 1998, and 2004 for the HRS, WB, and EBB cohorts, respectively). The sample includes husbands and wives who are observed through the entire period of observation; total number of observations is 7324. All dollar amounts are in terms of real 2005 dollars.

Appendix Table C

Conditional expected vs realized inheritance amounts, sample of inheritance recipients.

Cond. exp. inh. amount	Realized inheritance amounts					Total obs.
	>0 to 10k	>10k to 50k	>50k to 250k	>250k to 1M	>1M	
Zero	17.0	40.7	33.8	6.6	1.9	364
>0 to 10k	32.4	47.1	20.6	0.0	0.0	34
>10k to 50k	14.6	49.6	28.5	6.6	0.7	137
>50k to 250k	2.8	27.5	56.0	10.9	2.8	284
>250k to 1M	1.5	19.4	37.3	29.9	11.9	67
>1M	25.0	25.0	25.0	0.0	25.0	4
Missing	16.8	51.4	27.1	2.9	1.8	280
Total						1170

Notes: Row percentages are shown in each cell. Realized inheritance receipts are cumulative from the first period of observation through 2008. For the HRS, WB, and EBB cohorts the first period of observation is 1996, 2000, and 2006, respectively. The probability of receiving inheritance is the expectation as of the wave prior to the period of observation (i.e. 1994, 1998, and 2004 for the HRS, WB, and EBB cohorts, respectively). The sample includes husbands and wives who are observed through the entire period of observation; all dollar amounts are in terms of real 2005 dollars.

Appendix Table D

Descriptive statistics

Variable	Male		Female	
	Mean	Std. dev.	Mean	Std. dev.
<i>Dependent variable</i>				
Labor force participation	0.608	(0.488)	0.499	(0.500)
<i>Wealth and pensions</i>				
Net worth	5.212	(17.335)	5.212	(17.335)
DC account balance missing	0.088	(0.283)	0.068	(0.252)
DB pension wealth missing	0.116	(0.320)	0.082	(0.275)
DB pension wealth for exit at survey date	1.00	(2.341)	0.459	(1.678)
Gain in DB pension wealth for exit at 65	−0.099	(0.782)	−0.034	(0.550)
DC account balance	0.234	(2.225)	0.056	(0.333)
SS wealth for exit and claiming at survey date	1.285	(0.915)	0.944	(0.674)
Gain in SS wealth for exit and claiming at 65	0.259	(0.248)	0.254	(0.219)
SS wealth for entry to SSDI at survey date	1.278	(1.257)	1.2	(1.158)
PDV of lifetime earnings at 65	9.164	(7.332)	3.523	(4.443)
<i>Other control variables</i>				
Education = High school	0.339	(0.473)	0.414	(0.493)
Education = Some college	0.183	(0.387)	0.219	(0.414)
Education = College graduate	0.282	(0.450)	0.193	(0.395)
Hispanic	0.095	(0.293)	0.096	(0.294)
Black	0.097	(0.296)	0.094	(0.292)
Other race	0.044	(0.204)	0.044	(0.205)
Parent died since last interview	0.072	(0.259)	0.080	(0.271)
Employer retiree health insurance indicator	0.235	(0.424)	0.143	(0.350)
Retiree health insurance missing	0.279	(0.449)	0.557	(0.497)
Employer-provided health insurance indicator	0.498	(0.500)	0.290	(0.454)
Employer health insurance missing	0.021	(0.145)	0.024	(0.153)
Health good	0.317	(0.465)	0.300	(0.458)
Health fair poor	0.209	(0.407)	0.188	(0.390)
Health improved since last interview	0.097	(0.296)	0.103	(0.305)
Health worsened since last interview	0.199	(0.399)	0.193	(0.394)
Self employed	0.171	(0.377)	0.084	(0.277)
Age	62.3	(4.7)	59.7	(5.0)
Age = 62–64	0.200	(0.400)	0.198	(0.398)
Age = 65 +	0.369	(0.483)	0.193	(0.395)
Probability live to age 75	43.6	(37.6)	58.6	(33.9)

Notes: Total sample size is 27,400 person-wave observations (13,700 observations for each gender). See Table 1 for descriptive statistics on the inheritance variables. All dollar amounts are in units of \$100,000, deflated to year 2005.

Appendix Table E

Other coefficient estimates for specification in Table 3.

	Men		Women	
Inheritance receipt self (1 lag)	−0.033	(0.020)	−0.009	(0.021)
Inheritance receipt spouse (1 lag)	−0.014	(0.021)	−0.027	(0.022)
Inheritance receipt unknown (1 lag)	0.002	(0.023)	0.004	(0.020)
Inheritance receipt self (2 lag)	−0.008	(0.018)	−0.044	(0.020)**
Inheritance receipt spouse (2 lag)	0.024	(0.021)	0.022	(0.018)
Inheritance receipt unknown (2 lag)	0.007	(0.025)	0.010	(0.025)
Net worth (1 lag)	0.000	(0.000)	0.000	(0.000)
Net worth (2 lag)	0.000	(0.000)	0.000	(0.000)
DB pension wealth missing	0.028	(0.014)*	0.016	(0.016)
DB pension wealth of spouse missing	−0.016	(0.014)	0.001	(0.013)
DC account balance missing	−0.023	(0.012)*	−0.016	(0.014)
DC account balance of spouse missing	−0.012	(0.014)	−0.014	(0.012)
SS wealth of spouse missing	−0.023	(0.012)*	−0.017	(0.014)
DB pension wealth for exit at survey date	−0.003	(0.002)**	−0.001	(0.002)
Gain in DB pension wealth for exit at 65	−0.004	(0.004)	0.009	(0.006)
DC account balance	−0.001	(0.001)	−0.006	(0.008)
SS wealth for exit and claiming at survey date	−0.012	(0.010)	0.007	(0.013)
Gain in SS wealth for exit and claiming at 65	0.033	(0.018)*	0.031	(0.021)
SS wealth for entry to SSDI at survey date	0.005	(0.005)	−0.005	(0.004)
PDV of lifetime earnings at 65	0.000	(0.001)	0.002	(0.001)
DC pension indicator	0.041	(0.009)***	0.030	(0.011)***
DB pension indicator	−0.064	(0.012)***	−0.016	(0.012)
DB pension coverage indicator spouse	0.004	(0.012)	−0.011	(0.010)
DC pension coverage indicator spouse	0.012	(0.011)	0.009	(0.009)
DB pension wealth of spouse for exit at survey date	−0.003	(0.002)**	−0.001	(0.002)
Gain in DB pension wealth of spouse for exit at 65	−0.004	(0.004)	0.009	(0.006)
DC account balance of spouse	−0.001	(0.001)	−0.006	(0.008)
SS wealth of spouse	−0.012	(0.010)	0.007	(0.013)
Gain in SS wealth of spouse for exit and claiming at 65	0.035	(0.020)*	−0.009	(0.018)
SS wealth for entry of spouse to SSDI at survey date	−0.002	(0.005)	0.004	(0.004)
PDV of spouse's lifetime earnings at 65	−0.001	(0.001)	−0.001	(0.001)
Employer retiree health insurance indicator	−0.046	(0.010)***	−0.057	(0.011)***
Retiree health insurance missing	−0.025	(0.011)**	0.000	(0.015)
Employer-provided health insurance indicator	0.004	(0.009)	0.062	(0.014)***
Employer health insurance missing	0.018	(0.023)	0.020	(0.024)
Employer-provided health insurance indicator for spouse	−0.009	(0.007)	−0.014	(0.008)*
Employer health insurance missing for spouse	−0.007	(0.021)	−0.025	(0.023)
Education = High school	0.009	(0.009)	0.014	(0.009)
Education = Some college	0.014	(0.010)	0.017	(0.010)*
Education = College graduate	0.038	(0.010)***	0.024	(0.012)**
Hispanic	0.011	(0.011)	−0.003	(0.011)
Black	−0.007	(0.010)	0.014	(0.011)
Other race	0.013	(0.013)	0.009	(0.014)
Parent died since last interview	−0.008	(0.011)	−0.003	(0.011)
Health good	−0.013	(0.007)*	−0.007	(0.007)
Health fair poor	−0.094	(0.010)***	−0.071	(0.009)***
Health improved since last interview	−0.012	(0.010)	−0.024	(0.009)***
Health worsened since last interview	−0.038	(0.008)***	−0.036	(0.008)***
Self employed	0.050	(0.010)***	0.002	(0.015)
Self employed, spouse	0.008	(0.011)	0.011	(0.010)
Age	−0.006	(0.002)***	−0.009	(0.001)***
Age = 62–64	−0.087	(0.012)***	−0.058	(0.011)***
Age = 65+	−0.107	(0.019)***	−0.031	(0.021)
Probability live to age 75	−0.00003	(0.00010)	−0.00004	(0.00010)
Probability live to age 75, Spouse	−0.00002	(0.00009)	−0.00004	(0.000094)
R squared (sample size)	0.553 (13,724)		0.572 (13,724)	

Notes: The specification also includes Census Region dummies and survey year dummies. All dollar amounts are in units of \$100,000 (deflated to 2005). See Tables 2 and 3 for additional notes.

Appendix Table F


Estimated effect of amount inherited on LFP, alternative samples.

Sample cutoff	Nbr. obs.	Men			Women		
		Self	Spouse	Unknown	Self	Spouse	Unknown
None	13,724	0.010 (0.007)	−0.015 (0.009)	−0.019 (0.018)	−0.014 (0.009)	−0.002 (0.007)	0.011 (0.010)
99th percentile	13,718	−0.001 (0.006)	−0.021* (0.011)	−0.086* (0.045)	−0.021* (0.011)	−0.002 (0.009)	0.047* (0.027)
98th percentile	13,710	−0.006 (0.010)	−0.025* (0.015)	−0.086* (0.045)	−0.031** (0.013)	−0.005 (0.014)	0.047* (0.027)
97th percentile	13,702	−0.019 (0.015)	−0.018 (0.014)	−0.085* (0.048)	−0.036*** (0.014)	−0.003 (0.017)	0.045 (0.034)

Appendix Table F (continued)

Sample cutoff	Nbr. obs.	Men			Women		
		Self	Spouse	Unknown	Self	Spouse	Unknown
96th percentile	13,696	−0.024 (0.016)	0.003 (0.011)	−0.098** (0.050)	−0.040** (0.017)	−0.001 (0.019)	0.044 (0.037)
95th percentile	13,687	−0.034 (0.021)	−0.003 (0.014)	−0.117** (0.058)	−0.031* (0.019)	−0.007 (0.025)	0.046 (0.043)

Notes: The first column of each row indicates the percentile cutoff value for inheritance receipt; observations with inheritances greater than the value associated with the indicated percentile are trimmed from the sample. * (**) [***] indicates that the coefficient estimate is significantly different from zero at the 10% (5%) [1%] level.



Appendix Fig. A. Labor force participation rates of husbands and wives, by age. Notes: 13,724 household-wave observations. LFP rates at ages 45–49 are suppressed from the figure due to small sample sizes.

References

- Altonji, J., Elder, T., Taber, C., 2005. Selection on observed and unobserved variables: assessing the effectiveness of catholic schools. *J. Polit. Econ.* 113 (1), 151–184.
- Bobonis, G.J., 2009. Is the allocation of resources within the household efficient? New evidence from a randomized experiment. *J. Polit. Econ.* 117 (3), 453–503.
- Brown, J., Coile, C., Weisbenner, S., 2006. The effect of inheritance receipt on retirement. NBER Working Paper 12386.
- Brown, J., Coile, C., Weisbenner, S., 2010. The effect of inheritance receipt on retirement. *Rev. Econ. Stat.* 92 (2), 425–434.
- Browning, M., Chiappori, P.-A., Weiss, Y., 2014. *Economics of the Family*. Cambridge University Press, New York.
- Bruze, G., Svarer, M., Weiss, Y., 2015. The dynamics of marriage and divorce. *J. Labor Econ.* 33 (1), 123–170 January.
- Casanova, M., 2010. Happy together: a structural model of couples' joint retirement choices. Working Paper. Department of Economics, UCLA November 2.
- Cesarini, D., Lindqvist, E., Notodowigdo, M.J., Ostling, R., 2015. The effect of wealth on household labor supply: evidence from Swedish lotteries. NBER Working Paper.
- Chiappori, P.-A., Fortin, B., Lacroix, G., 2002. Marriage market, divorce legislation, and household labor supply. *J. Polit. Econ.* 110 (1), 37–72.

- Duflo, E., 2003. Grandmothers and granddaughters: old age pension and intra-household allocation in South Africa. *World Bank Econ. Rev.* 17 (1), 1–25.
- Duflo, E., Udry, C., 2004. Intrahousehold resource allocation in Cote D'Ivoire: social norms, separate accounts, and consumption choices. NBER Working Paper 10498 May.
- Friedberg, L., Webb, A., 2006. Determinants and consequences of bargaining power in households. NBER Working Paper 12367 July.
- Gallipoli, G., Turner, L., 2013. Social security, endogenous retirement and intrahousehold cooperation. Working Paper. University of British Columbia September 16.
- Gemici, A., 2011. Family migration and labor market outcomes. Working Paper, NYU.
- Gouskova, E., 2013. Inheritance Reporting in the Health and Retirement Study Data: Evidence of Forward Telescoping. <http://ssrn.com/abstract=2267507>.
- Gustman, A.L., Steinmeier, T.L., 2009. Integrating retirement models. NBER Working Paper 15607 December.
- Holtz-Eakin, D., Joulfaian, D., Rosen, H.S., 1993. The Carnegie conjecture: some empirical evidence. *Q. J. Econ.* 108 (2), 413–435.
- Hurd, M., McFadden, D., Gan, L., 1998. Subjective survival curves and life cycle behavior. In: Wise, D.A. (Ed.), *Inquiries in the Economics of Aging*. University of Chicago Press for the NBER, pp. 259–305.
- Imbens, G.W., Rubin, D.B., Sacerdote, B.I., 2001. Estimating the effect of unearned income on labor earnings, savings, and consumption: evidence from a survey of lottery players. *Am. Econ. Rev.* 91 (4), 778–794.
- Joulfaian, D., Wilhelm, M.O., 1994. Inheritance and labor supply. *J. Hum. Resour.* 29 (4), 1205–1234.
- Laitner, J., Sonnega, A., 2010. Intergenerational transfers in the health and retirement study data. Working Paper 2010-238. Michigan Retirement Research Center.
- Lich-Tyler, Stephen. No date. "Negotiations and love songs: the dynamics of bargained household decisions," Working Paper, University of Michigan, Ann Arbor.
- Lise, J., Yamada, K., 2014. Household Sharing and Commitment: Evidence from Panel Data on Individual Expenditures and Time Use. Singapore Management University.
- Lundberg, S.J., Pollak, R.A., Wales, T.J., 1997. Do husbands and wives pool their resources? Evidence from the U.K. child benefit. *J. Hum. Resour.* 32 (3), 463–480.
- Lundberg, S., Startz, R., Stillman, S., 2003. The retirement-consumption puzzle: a marital bargaining approach. *J. Public Econ.* 87, 1199–1218.
- Maestas, N., 2001. Labor, love & leisure: complementarity and the timing of retirement by working couples. Working Paper. University of California, Berkeley December.
- Marcel, A., Marimon, R., 2011. Recursive contracts. Max Weber Program Working Paper 2011/3. European University Institute, Florence, Italy.
- Martin, P.W., 2012. The case for reforming the program's spouse benefits while "Saving Social Security" (2012). Cornell Law Faculty Working Papers. Paper 101 http://scholarship.law.cornell.edu/clsoops_papers/101.
- Mazzocco, M., 2007. Household intertemporal behavior: a collective characterization and a test of commitment. *Rev. Econ. Stud.* 74, 857–895.
- Mazzocco, M., Ruiz, C., Yamaguchi, S., 2007. Labor Supply, Wealth Dynamics, and Marriage Decisions. Working Paper UCLA, October.
- Persson, P., 2014. Social insurance and the marriage market. Manuscript. Stanford University October.
- van der Klaauw, W., Wolpin, K., 2008. Social security and the retirement and savings behavior of low-income households. *J. Econ.* 45 (1–2), 21–42.
- Voena, A., 2015. Yours, mine, and ours: do divorce laws affect the intertemporal behavior of married couples? *Am. Econ. Rev.* 105 (8) 2295–32.